REPORT DOCUMENTATION PAGE

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE March 15, 1995	3. REPORT TYPE AND DATES COVERED

4. TITLE AND SUBTITLE
Intelligent Graphic Interfaces for Displaying Large Amounts of Information

5. FUNDING NUMBERS
DAAH04-93-G-0507

6. AUTHOR(S)
Steven F. Roth

DTIC SELECTED
OCT 2 4 1995
S G D

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Carnegie Mellon University
The Robotics Institute
5000 Forbes Ave
Pittsburgh, PA 15213-3890

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
U.S. Army Research Office
P.O. Box 12211
Research Triangle Park, NC 27709-2211

10. SPONSORING/MONITORING AGENCY REPORT NUMBER
ARO 32140.3-MA

11. SUPPLEMENTARY NOTES
The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy, or decision, unless so designated by other documentation.

12a. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited.

12b. DISTRIBUTION CODE

13. ABSTRACT (Maximum 200 words)

Research was conducted on intelligent interfaces for exploring and visualizing large amounts of diverse information for supporting the decision-making process. This research combined technical approaches for automatically generating visualizations of information with an approach to human-computer interaction that developed interactive techniques for performing data manipulation tasks with these visualizations. The important results of this research project include:
- the development of multiple interactive techniques for exploring information in visualizations,
- the coordination of these techniques within a single environment so they can be used in a concerted way, and
- the integration of these techniques within the SAGE data graphic design system so that they can be applied in a consistent manner across potentially hundreds of different graphics that vary along many dimensions.

1995 1023 006

DTIC QUALITY INSPECTED 1

14. SUBJECT TERMS
Data Visualization, Graphics, Data Exploration

15. NUMBER OF PAGES
4

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UL

Intelligent Graphic Interfaces for Displaying Large Amounts of Information

Final Report
Dr. Steven F. Roth

March 15, 1995

U.S. Army Research Office
Contract Number DAAH04-93-G-0507

Carnegie-Mellon University
5000 Forbes Avenue
Pittsburgh PA 15213

Approved for public release;
Distribution Unlimited

Accesion For		
NTIS CRA&I		☑
DTIC TAB		☐
Unannounced		☐
Justification		
By		
Distribution /		
Availability Codes		
Dist	Avail and / or Special	
A-1		

Abstract

Research was conducted on intelligent interfaces for exploring and visualizing large amounts of diverse information for supporting the decision-making process. This research combined technical approaches for automatically generating visualizations of information with an approach to human-computer interaction that developed interactive techniques for performing data manipulation tasks with these visualizations. The important results of this research project include:

- the development of multiple interactive techniques for exploring information in visualizations,
- the coordination of these techniques within a single environment so they can be used in a concerted way, and
- the integration of these techniques within the SAGE data graphic design system so that they can be applied in a consistent manner across potentially hundreds of different graphics that vary along many dimensions.

Problems Studied & Summary of Important Results

There have been several goals of our research during the course of this project. Our long term goal is to develop environments in which people can explore and analyze large amounts of diverse information. These environments must have usable mechanisms for creating effective visual representations of data which support the tasks they are performing. Our approach to this problem has been to integrate knowledge-based techniques for graphic design with interactive techniques for exploring and manipulating quantitative and relational data. As we discussed in previous reviews, there were several basic research problems that needed to be addressed. The focus of the research we have worked on within this contract has been concerned with exploring interactive techniques for manipulating data within automatically generated displays produced by our knowledge-based data graphic presentation system (SAGE).

In the progression of our work on SAGE, an automatic knowledge-based data visualization tool, we have developed an architecture which enables a number of interactive techniques to be used to control all SAGE-generated visualizations. With these techniques, every picture produced by SAGE becomes a flexible interface for filtering and understanding data.

Visualizations vary dramatically in their complexity. We've focused our work on what we call integrative graphics: namely those that show numerous properties of the underlying data in an integrated fashion. In the example, you can scan for the activity with the shortest duration by searching for the smallest interval bar. From this bar it's easy to determine the cost of the activity (by looking at the size of the clustered circle) or the manager (by looking to the right on the associated table). In an integrative visualization identifying a relevant assertion facilitates the lookup of other related assertions.

Graphical integration can be achieved using combinations of several different graphical techniques: showing data with multiple parameters of graphical objects, showing data in

multiple graphical objects by clustering them in the same space (frame of reference), or by showing them in separate spaces, which can be aligned or otherwise integrated.

Many interactive interfaces have been studied for the task of visual information seeking. These interfaces provide users with the capability to:

filter data, through the progressive refinement of queries [dynamic query, infocrystal]

see views of their data in multiple visualizations, and highlight items coordinated across these views [painting, brushing, parallel coordinates]

initially see data at a high level of abstraction, and drill-down along orthogonal properties of the dataset to show progressively more detail [spreadsheets, PowerPlay]

manipulate representations at a high level of abstraction, but see detail when desired [Aggregate Manipulator, table lens, magic lens]

Many of these techniques have been implemented in coordination with one or more visualizations, allowing users to more quickly evaluate the results of their operations, thereby enabling faster feedback.

Our study of these techniques revealed that an integration of several such techniques would greatly improve the usability of a visual information seeking tool. We wanted to provide our users with access to these and other new interactive techniques on a wide variety of integrative displays. Unfortunately, many of these techniques were only implemented with very simple, non-integrative visualizations. Most expressed no systematic approach of how the interface operations were to tie into a visualization.

We performed an analysis of the various techniques in the context of our understanding and representation of the syntax and semantics of graphics. We developed a framework for describing interactive techniques, and developed a dataflow architecture built on top of SAGE's design component enabling the design of interactive visualizations.

Our initial work focused on several capabilities:

painting - Users select graphical elements in one display by either selection a region or by clicking on individual element. Each selected element and all other views of data underlying these elements are painted (e.g. colored red).

dynamic query sliders - Sliders whose values range over a specified domain of the data are shown in the visualization. When users manipulate the sliders all graphical elements expressing data that falls outside the specified range are hidden.

composition/decomposition - Using the aggregate manipulator interface or my directly performing operations on graphical elements users control the level of abstraction the of the presented data. For example, the user can begin with a single graphical element representing all the data in the database, and then, using the decomposition operation repeatedly partition the points in this abstraction by a particular attribute, thereby, drilling down to an interesting subset. Similarly, using the composition operation the user can select elements and roll them up into a new element.

multiple coordinated displays - Users can form a set of coordinated displays. In such a set, an interface action in one is reflected in all the displays. For example, when a group of items are painted in one display, the elements corresponding to the selected data

are also painted in all other coordinated displays. In this context, painting provides a means to integrate separate displays.

The integration of these four techniques demonstrates the generality of our system, and forms a minimal useful set of data manipulation techniques to enable SAGE's use as a testbed for data exploration tasks. Note that our architecture can be used to easily implement a large number of interactive techniques.

The second major advantage to the generality of the flexible design language is that we were able to build an instantiation and rendering module which takes a language plus a data set as input and generates a complete display as output. Thus, any graphic that can be constructed using the design language can be turned into an actual display. Furthermore, the actually mapping process between specific data objects and attributes and graphic properties is automated because of the expressiveness criteria stored with each language. Thus, a user need only specify the data and language to be used and the system uses its knowledge to map attributes to graphic properties. Finally, the languages are such that they can be composed, so that two graphics can be merged (e.g. aligning multiple charts, integrating a network with a map, merging gauges and textual displays with charts or maps).

The second major area of research has been interactive methods for supporting exploration and manipulation large amounts of multi-dimensional, heterogeneous information. We identified a set of interactive capabilities required to perform many logistical tasks:
• Identifying subsets of information relevant to current tasks (e.g., to perform search, to partition available information, to define ranges of information, etc.)
• Controlling the level of detail with which information is displayed, including the ability to specify computations which are necessary to aggregate, abstract, and summarize information (conversely, to decompose or increase the level of detail with which information is expressed)
• Specifying the focus of attention: aspects of the information that are important to display (e.g., specify the attributes, features, or characteristics of information that are relevant to the current task)

List of Publications

"A Framework for Knowledge-based Interactive Data Exploration", Goldstein, J., Roth, S.F., Kolojejchick, J., Mattis, J., Journal of Visual Languages and Computing, December, 1994.

"Using Aggregation and Dynamic Queries to Explore Large Data Sets", Goldstein, J. and Roth, S.F., Proceedings SIGCHI'94 Human Factors in Computing Systems, Boston, MA, ACM, April 1994.

Participating Scientific Personnel

Steven F. Roth, PhD
John Kolojejchick